

Re-Ranking Browse for Everyone — Including Customers We Know Nothing About

Launching a browse re-ranker ML model that personalizes product ordering for both identified and anonymous users, driving tens of millions of dollars in incremental annual revenue.

Context

Browse pages — category and collection listings — are among the highest-traffic surfaces in e-commerce, yet they were ordered by static, one-size-fits-all logic. Every customer landing on “Sofas” saw the same wall of products, whether they were a returning customer with rich history or a first-time visitor from a search ad.

The Problem

The obvious ML answer — personalize using behavioral history — only works for identified users. A large share of revenue-driving traffic is anonymous: no login, thin or no history. Most personalization programs quietly ignore these users. I made the product call that we would not: a re-ranker that only improved the experience for logged-in customers would cap its own impact and widen the experience gap.

Approach & Key Decisions

- **One model, graceful signal degradation.** We designed the re-ranker to consume a spectrum of signals — long-term preferences for identified users, in-session behavior (clicks, dwell, filters) for anonymous ones, and contextual signals (entry point, device, geography) for everyone. The model degrades gracefully as signal thins, rather than requiring a separate cold-start system.
- **Session as the unit of intent.** For anonymous users, the key reframe was treating the current session as a rich signal source in its own right. A few clicks into a browse session, intent is often clearer than a year of stale purchase history.
- **Relevance with guardrails.** Pure engagement optimization can degrade assortment diversity and strand long-tail suppliers. We built diversity and business-rule constraints into the ranking layer and monitored them as first-class metrics, not afterthoughts.
- **Experimentation rigor.** Browse touches merchandising, ads, and supplier outcomes. I drove a staged rollout with pre-registered success criteria, holdout groups for long-term effects, and per-surface analysis so wins on one page type couldn't mask losses on another.

Impact

Revenue	
	Tens of millions of dollars in estimated incremental annual revenue from improved browse relevance (directional; exact figures confidential).

Coverage	Personalization extended to anonymous traffic — the majority of visitors — not just logged-in customers.
Platform	The ranking framework and signal pipeline became the foundation for personalization on additional surfaces.

What I'd Tell Other PMs

“Cold start isn't an ML edge case — it's most of your traffic. Design for the customer you know nothing about, and the rest gets easier.”

The strategic decision (who gets personalization) mattered more than any modeling choice. Insisting that anonymous users were in scope forced better architecture, honest metrics, and a much larger prize.

Sneha Chitte · Senior Product Manager – ML & Personalization · Greater Boston, MA · Metrics are directional and anonymized to protect confidential information.